

Benchmarking Algorithmic Bias in Face Recognition: An Experimental Approach Using Synthetic Faces and Human Evaluation

Hao Liang
Rice University
hl106@rice.edu

Pietro Perona
California Institute of Technology and AWS
perona@caltech.edu, peronapp@amazon.com

Guha Balakrishnan
Rice University
guha@rice.edu

Abstract

We propose an experimental method for measuring bias in face recognition systems. Existing methods to measure bias depend on benchmark datasets that are collected in the wild and annotated for protected (e.g., race, gender) and unprotected (e.g., pose, lighting) attributes. Such observational datasets only permit correlational conclusions, e.g., “Algorithm A’s accuracy is different on female and male faces in dataset X.” By contrast, experimental methods manipulate attributes individually and thus permit causal conclusions, e.g., “Algorithm A’s accuracy is affected by gender and skin color.”

Our method is based on generating synthetic faces using a neural face generator, where each attribute of interest is modified independently while leaving all other attributes constant. Human observers crucially provide the ground truth on perceptual identity similarity between synthetic image pairs. We validate our method quantitatively by evaluating race and gender biases of three research-grade face recognition models. Our synthetic pipeline reveals that for these algorithms, accuracy is lower for Black and East Asian population subgroups. Our method can also quantify how perceptual changes in attributes affect face identity distances reported by these models. Our large synthetic dataset, consisting of 48,000 synthetic face image pairs (10,200 unique synthetic faces) and 555,000 human annotations (individual attributes and pairwise identity comparisons) is available to researchers in this important area.

1. Introduction

Face recognition technology has applications in consumer media, information security, access control, law enforcement, and surveillance systems. The one-to-one task is to determine whether a pair of faces share the same identity (“face verification”). The one-to-many task determines whether a face image shares an identity with one or more from a database of face images from known individuals

(“face identification”). Face recognition systems implemented with deep neural networks today achieve impressive accuracies [54, 13, 33, 43] and outperform even expert face analysts [38]. Nevertheless, it is important to detect and measure possible algorithmic *biases*, i.e., systematic accuracy differences, especially across protected demographic attributes like age, race and gender [10, 22, 20], in order to maintain fair treatment in sensitive applications. For this reason, the National Institute of Standards and Technology (NIST) measures bias in commercial face recognition models [17], in particular by comparing their False Match Rate (FMR) and False Non Match Rate (FNMR) values across different demographic subgroups at a particular decision threshold (sweeping this threshold yields FNMR vs. FMR “curves”).

The first step in measuring bias of face recognition systems is, currently, to collect a large benchmarking dataset containing a set of diverse faces, where each is photographed multiple times under different conditions. An algorithm’s error rate across subgroups specified by different protected attribute combinations (e.g., different race and gender groups) can then be measured.

Unfortunately, sampling a good test dataset is almost impossible. First, each protected intersectional group (a specific combination of attribute values) must contain a sufficiently large number of individuals, which is not easy when sampling faces in the wild. Second, in order to estimate the *causal* effect of these protected attributes on algorithm bias, it is crucial to ensure that the joint distribution of unprotected attributes like age, lighting, pose, etc. is roughly equal across protected attribute groups – otherwise biases in the dataset will be misinterpreted as biases in the algorithm. It is practically impossible to do this with *observational* data, i.e., data that we have no power to construct or intervene on ourselves, and all available datasets fall short of this criterion. Third, privacy concerns with collecting human data, and the cost and accuracy of identity annotation, make dataset construction for face recognition benchmarking very expensive.

We address this challenge by developing an approach to

construct *synthetic* face recognition benchmarking datasets using a combination of modern face generation models and human annotators. We assume access to a pretrained face generator with a latent space, such as any of the popular public models trained in a generative adversarial network (GAN) framework [16, 32, 24, 9, 25]. We traverse the generator’s latent space to first construct many “identities” (actually, “pseudo-identities” since there is no real person behind the face images) spanning different protected attributes (race and gender in our experiment). For each identity, we vary several unprotected attributes (pose, age, expression and lighting) to construct face sets that ideally depict the same person in different settings. We vary the unprotected attributes in the same way for all faces to ensure a balanced dataset. A challenge in using synthetic faces is we do not have a ground truth identity label per face. Such labels are needed to evaluate the accuracy of our algorithms. We replace the ground truth with consensus from human annotations, which we call the “Human Consensus Identity Confidence” (HCIC). We use the HCICs along with the synthetic images to benchmark face recognition systems for bias.

Our approach is fast, practical, inexpensive, and virtually eliminates privacy concerns in data collection. The closest related work uses synthetic faces to benchmark face analysis systems [6], which classify attributes like gender and expression for a single face. We build on their ideas, but address the unique challenge of obtaining ground truth annotations for identity comparisons *between face pairs*. Our work is the first to present an annotation pipeline for verifying a synthetic face recognition benchmarking dataset, and to demonstrate that a synthetic approach can be reliable for causal face recognition benchmarking.

Using our synthetic dataset consisting of 48,000 synthetic face image pairs (10,200 unique synthetic faces) and 555,000 human annotations (individual attributes and pairwise identity comparisons), we computed face identity distances reported by three public face recognition models. We used the HCICs to assign a ground truth to each pair, and computed False Non Match Rate (FNMR) vs. False Match Rate (FMR) curves for different attributes and demographic groups. Our results show that these algorithms have lowest error rates on White (a shorthand for European-looking, or Caucasian) groups. For each model, we also report the expected change to predicted face identity distance with respect to changes to each unprotected attribute.

We make three contributions: (a) A method to *experimentally* estimate bias in face verification algorithms, and thus estimate the causal connection between attributes and bias, eliminating confusion between test set and algorithmic bias. (b) An empirical evaluation of our method, where we discover lower accuracy for Black and East Asian faces compared to Caucasian faces in three popular academic-grade face verification algorithms. (c) A large dataset of

synthetic faces with human-collected ground truth.

2. Related work

Face recognition. Face recognition [54, 13, 33, 43] is one of the most successful applications of image analysis and understanding, and has been applied to many areas including information security, access control, law enforcement, and surveillance systems. State-of-the-art models all use deep neural networks [13, 33, 43, 50], and typically work by embedding each face into a latent space, and assigning an identity distance between a pair of faces by computing a distance (typically cosine distance) in that latent space. These models are trained over large-scale face datasets [37, 34, 4, 14], most of which contain thousands of identities and multiple images per identity.

Fairness in computer vision. There is an increasing focus on measuring and mitigating biases of computer vision models and datasets [17, 56, 41, 40, 15, 28, 45]. Biases may be measured with a number of metrics [47, 18, 11] to quantify performance differences of algorithms across population subgroups. Face recognition and analysis systems are often under the most scrutiny due to their sensitive nature [7, 6, 23, 56]. A typical approach to measure bias of a face analysis algorithm is to collect a real dataset of faces with appropriate labels of protected attributes like race and gender, and compute model error across intersection subgroups of these attributes [7, 2, 1, 12, 30, 35]. Recent work argues that measuring bias on datasets that are collected in the wild risks confusing dataset bias with algorithmic bias – the fix is using an experimental approach using synthetic face generators [6]. As an added bonus using synthetic images lowers the cost and privacy concerns of data collection. We build on that study in this work, by extending their ideas to benchmark face recognition models.

Bias mitigation. Methods for mitigating bias mainly focus on two elements: the training dataset and the model design. A dataset has sampling biases if its joint distribution of attributes is far from random. For example, the CelebA face dataset is known to have significant sampling biases, such as a higher proportion of female with young ages compared to men [6]. A model trained on a biased dataset may inherit biases [23, 52, 39], particularly for attribute subgroups that are underrepresented in the dataset. Conversely, bias can be mitigated if the algorithm is trained on a balanced dataset. Researchers are also building novel model designs to combat biases [53, 57, 26, 48, 49], typically by learning representations invariant to protected attributes or sampling in a balanced way during training.

Face generation. Several algorithms can produce near photorealistic synthetic faces: Variational Autoencoders (VAE) [29, 46, 31], GANs [16, 32, 24, 9], and diffusion models [21, 36]. We use StyleGAN [24, 25, 9]. A useful property of a GAN is that it generates images from




Figure 1. **Examples of prototypes image spanning protected attributes of race and gender.** In the first step of our framework (see Sec. 3.1), we generate random “seed” images by randomly sampling in the latent space of our face generator. Starting from each seed, we perform latent space traversals (see Sec. 3.1) along latent dimensions correlated with race and gender to produce faces depicting different “demographic groups” (M=Male, F=Female, W=White, B=Black, A=East Asian). Each row in this figure shows the six prototypes generated from the same seed image. Each column shows members of the same *group* that were generated from different *seeds*. See Fig. 2 for the next step.

a low-dimensional latent space where it is often possible to find a one-dimensional subspace that alters a given scalar attribute like gender or age. The most common strategy is to train a linear classifier or regressor that predicts values of an attribute from the latent vector, and then traverse along the norm of the resulting model’s hyperplane to generate new images varying that attribute[44, 19, 42]. Higher-order methods that examine local latent space geometry and non-linear traversals have also been proposed[5].

3. Method

Our method consists of seven steps. First, generate *seed* face images (see Fig. 1) by sampling the latent space of the GAN. Second, generate *prototype* face images spanning all protected attribute combinations (race and gender in our experiments) by controlling features in the latent space of a face generator (see Fig. 1) (we use the abbreviations: {“WM”, “WF”, “BM”, “BF”, “AM”, “AF”} for {“White Male”, “White Female”, “Black Male”, “Black Female”, “East Asian Male”, “East Asian Female”}). Third, generate systematic modifications of each prototype image along unprotected semantic attributes (pose, age, expression, and lighting) (see Fig. 2). These form sets of faces that correspond to the same intended identity. Fourth, form pairs of images, both within and across identity sets. Fifth, give face pairs to human annotators to obtain perceived identity match scores which we call “Human Consensus Iden-

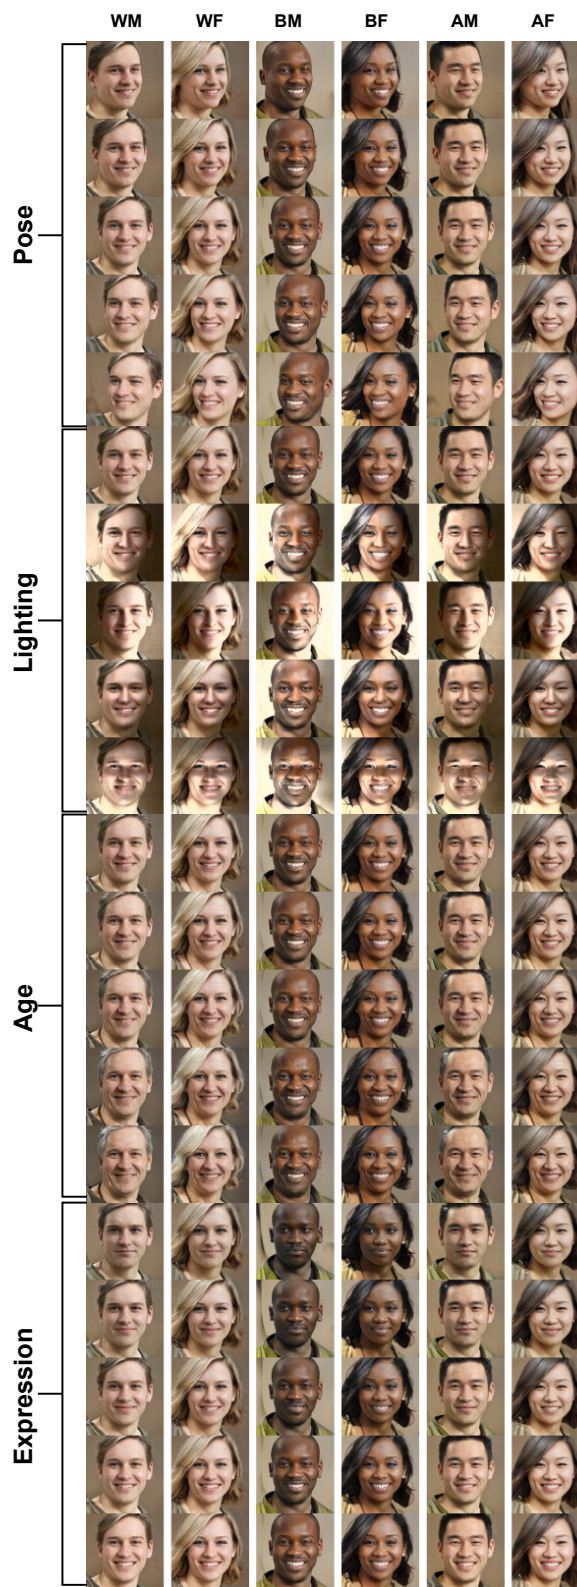


Figure 2. **Examples of modifying unprotected attributes.** In the second step of our framework (see Sec. 3.2), we modify each prototype face (see Sec. 3.1 and Fig. 1, 1st row) along attributes. All 120 images here come from one seed (S1 in Fig. 1). Each column (20 images) is generated from the same prototype.

tity Confidences” (HCICs). Sixth, ask human annotators to determine the attributes of single images and the similarity of image pairs. Seventh, feed the synthetic pairs to face recognition models, and use the HCICs to evaluate them for biases. We describe these steps in the next sections.

3.1. Generating images of different demographic groups along protected attributes

We assume a pretrained face generator that maps vectors $\mathbf{z} \sim p(z) \in R^d$ to images, where R^d is a low-dimensional latent space and p is a given probability density. In this work, we use a generator trained in a GAN framework. In particular, we use EG3D [9], a state-of-the-art generator with explicit control over geometry and pose.

We aim to generate a large set of “prototype” faces that depict all combinations of protected attribute intersectional subgroups. We consider the protected attributes of race (limited for simplicity to the East Asian, Black, Caucasian subgroups) and gender (similarly limited to the Male, and Female subgroups), resulting in 2 (genders) \times 3 (races) = 6 prototypes generated per seed. By using one seed to generate each set of 6 prototypes, our intention is to keep attributes other than race and gender approximately constant, so that the six seeds are *matched* vis-a-vis potential confound variables, which is useful for later statistical analysis.

Our method for modifying the seed images to produce prototypes is based on linear latent space traversals [6, 19, 42, 44]. We sample a large number of images from $p(z)$, and label race and gender of each using a public multi-task classifier [23]. This produces a training set $D = \{I^j, Z^j, L^j\}_{j=1}^N$, where I^j, Z^j, L^j denotes the face image, latent code, and labels of sample j respectively, and N is the total number of samples (100,000 in our experiments). Second, we train one support vector machine (SVM) to classify gender (assuming binary labels), and three one-vs-all SVMs for each of the three race subgroups. For a seed image, we generate a set of prototypes by starting at the seed’s position in latent space and moving along the normal directions of each attribute’s SVM hyperplane.

Not all generated images are realistic, and some of the seeds are similar in appearance to each other. To eliminate unrealistic images and quasi-duplicates, we filter the many random seeds into a smaller set that has diverse and realistic prototypes. We accomplish this by selecting 300 (out of 1000 initial) seeds using a human annotator judging the realism and agreement of the associated prototypes with respect to intended race and genders. Then, we further filter these 300 seeds into 100 using a max-min clustering algorithm (see Algorithm 1 in Supplementary) evaluated over the geometry of the faces as captured by a public 3D face mesh predictor [27]. This yields 600 total prototype images and a total of 10,200 images (17 per prototype, Fig. 2 and Sec. 3.2).

3.2. Modifying unprotected attributes

After creating prototype images of different demographic groups, we modify them according to various semantic attributes to build our dataset. We modify the following attributes: pose, lightning, age, and expression. We adopt different techniques when manipulating the attributes, as detailed in the following subsections. Starting from each prototype, we generate one image sequence per attribute depicting different gradations of change to that attribute. Examples of these sequences are shown in Fig. 2. In our experiments, we set all sequence lengths to 5.

Pose: Pose is explicitly modeled by the EG3D generator. For each prototype, we generate a sequence of images at angles of $\{-30^\circ, -15^\circ, 0^\circ, 15^\circ, 30^\circ\}$.

Lighting: We manipulate lighting using a public neural network person relighting model [55]. We used four lighting conditions, produced by a light source located $\{up, down, left, right\}$ with respect to the face. We set the power of the light source to 0.7, yielding a strong, but not overwhelming lighting effect.

Age and expression: We control these with the latent space traversal method mentioned in Sec. 3.1. We use pretrained models to assign age and expression labels to all face images in D , and train linear regressors to predict these labels from the latent vectors. To generate the sequences, we start at a prototype’s position in latent space and move along the normal vector of the regression hyperplane with a small step size until a certain score (0.8 for age, 0.9 for expression after normalizing the outputs of models to $[0, 1]$) is predicted by the model at a distance d from the prototype. Then, we take n steps from the prototype of size d/n along both directions. We set $n = 2$ in our experiments, resulting in $2n + 1(\text{prototype}) = 5$ images per image sequence.

3.3. Face segmentation and background removal

We further white-out all pixels not corresponding to the head (background, shoulders, neck) using a public face segmentation model [51]. This is for two reasons. First, some images have artifacts outside of the face (see Fig1 seed 4, for artifacts of fingers). Although most face recognition models first run face detection to crop regions outside of the face, we still remove the background for safety. Second, the background may inadvertently affect human annotator judgments. An example image before and after face segmentation and background removal is in Fig. 2 in Supplementary.

3.4. Generating face pairs

A face recognition test dataset consists of both ‘positive’ and ‘negative’ face pairs, where positive pairs depict the same identity (see Fig. 2) and negatives do not. We create positive pairs by taking each prototype face and pairing it with all other faces we generated starting from the

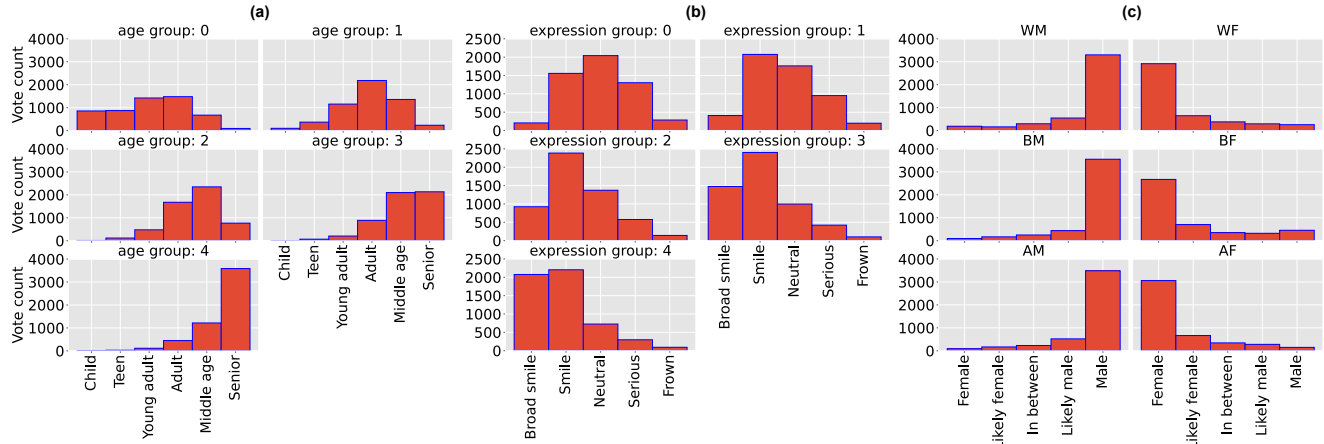


Figure 3. **Validation of attribute manipulation through single image annotation.** (a): Annotator scores for age, ranging from “child” to “senior”, on single synthetic images belonging to age groups 0 (youngest-looking) to 4 (oldest-looking) (Fig. 2). (b): Annotator scores for expression, ranging from “frown” to “broad smile”, each histogram is associated to a different expression group, from 0 (most frowning) to 4 (most smiling). (c): Annotator scores ranging from “female” to “male”. Histograms are shown for different demographic groups. The annotators’ perceptions track along with our manipulations and validate our method. See Fig. 3 in Supplementary for results on skin tone and uncanny.

prototype, varying one unprotected attribute at a time (see Fig. 2). We construct negative pairs by pairing each prototype image with faces from sequences produced by n other prototypes of the *same race and gender*, also varying one unprotected attribute at a time. In this way, a negative face pair will depict a change to a single attribute, along with a highly likely change in identity. Because the number of total possible negative pairs produced in this manner is large, we set $n = 3$ in our experiments. Therefore, our database consists of 12,000 positive pairs and 36,000 negative pairs.

3.5. Human annotation experiments

The attribute modifications we perform in Secs. 3.1 and 3.2 may not always work as intended. In particular, attributes can change at different rates in different sections of the generator’s latent space. More importantly, we have no ground truth measurements for face identity similarity between image pairs. To address both issues, we collect human annotations using *Amazon SageMaker Ground Truth*. For each image we collected 9 human annotations per attribute and average their responses to obtain one score.

Single image annotation: We give annotators **one** image from the synthetic dataset at a time, and ask them to annotate skin type, perceived gender, expression, age, and uncanniness (fakeness) using five-point scales.

Image pair annotation: To ensure that a pair of faces *does* belong to same/different person(s) from the perspective of a majority of humans, we collect pairwise ground truth annotations. For each pair of images 9 annotators are asked to choose from the following options regarding identity: {‘likely same’, ‘possibly same’, ‘not sure’, ‘possibly different’ and ‘likely different’}. A screenshot of the anno-

tation interface is in Fig. 1 in Supplementary.

4. Experiments

We evaluate our method on three public, research-grade face recognition models: a SpheroFace [33] model with SFNet-20 architecture trained on the VGGFace2 dataset [8] (Model 1), an Arcface [13] model with ResNet-34 architecture trained on the MS1MV3 dataset [3] (Model 2), and an Arcface [13] model with ResNet-34 architecture trained on the Glint360k dataset [14] (Model 3). All of the models were trained on large in-the-wild datasets and reached high accuracy on their respective test datasets. At inference time, for a pair of face images (I_A, I_B), each model returns feature vectors (f_A, f_B), and similarity between faces is quantified with cosine similarity: $(f_A \cdot f_B) / (\|f_A\| \|f_B\|)$.

Following the image generation pipeline introduced in Sec. 3.1 and Sec. 3.2, we build our synthetic dataset with 10,200 face images depicting 600 intended identities. Following the approach in Sec. 3.4, we use these images to produce 12,000 positive face pairs and 36,000 negative face pairs. We obtained 9 human annotations per image and attribute to quantify the unprotected attributes of age and expression (we did not annotate pose and lighting as our method of controlling them is much more precise), as well as 9 annotations per image pair to annotate identity distances. In total, we collected 123,000 annotations from 2,214 annotators (median=65, max=223 annotations each) for single image attribute annotations and 432,000 annotations from 1,905 annotators (median=264, max=402 annotations each) for image pair annotations.

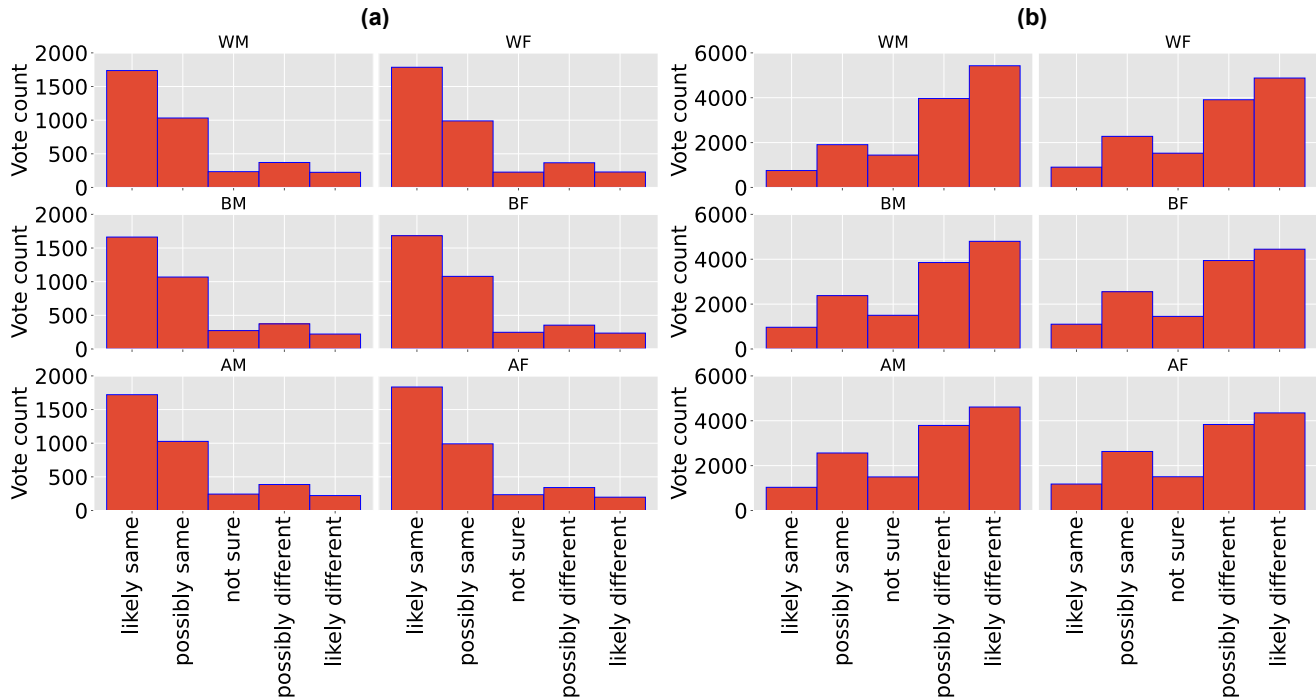


Figure 4. **Human annotation results on different pose image pairs.** (a): face images of different poses from the same seed (putative identity) and same demographic group. (b): face images of different poses from different seed (different putative identity) and same demographic group. Crucially, we do not find statistical differences in the similarity or difference of face pairs across different demographic groups. This is consistent across attributes, see Fig. 8 in Supplementary for results on other attributes.




Figure 5. **Identity similarity scores reported by a recognition model with respect to pose and demographic changes.** As described in Sec. 4.1.1, we feed face image pairs of different groups (see Fig. 1) to a public recognition model. Box plots show the median cosine distance (red line in the boxes) and the 15% – 85% confidence interval. The red dashed line at 0.6 indicates a reasonable threshold t to decide whether a pair of images are from the same person or not (the threshold trades off False Non Match Rate (FNMR) for False Match Rate (FMR) errors). The left plot corresponds to face pairs from the same seed and demographic group, i.e., nominally the same ‘individual.’ The second corresponds to face pairs from the same demographic group and different seed (i.e., different individuals with the same gender and similar racially-correlated attributes), and the last plot corresponds to face images from different demographic groups and seeds. It is clear that race and gender have a strong effect on the perception of identity. Results for other attributes are in Fig. 4 in Supplementary.

4.1. Annotation and image synthesis analysis

We first evaluate single image annotations for the attributes of age, skin color and gender in Fig. 3 (please see

figure caption for details). The results show that annotator perceptions agree well with our intended manipulations across all race and gender subgroups. We also asked annotators to label image realism (“uncanniness” in our surveys),

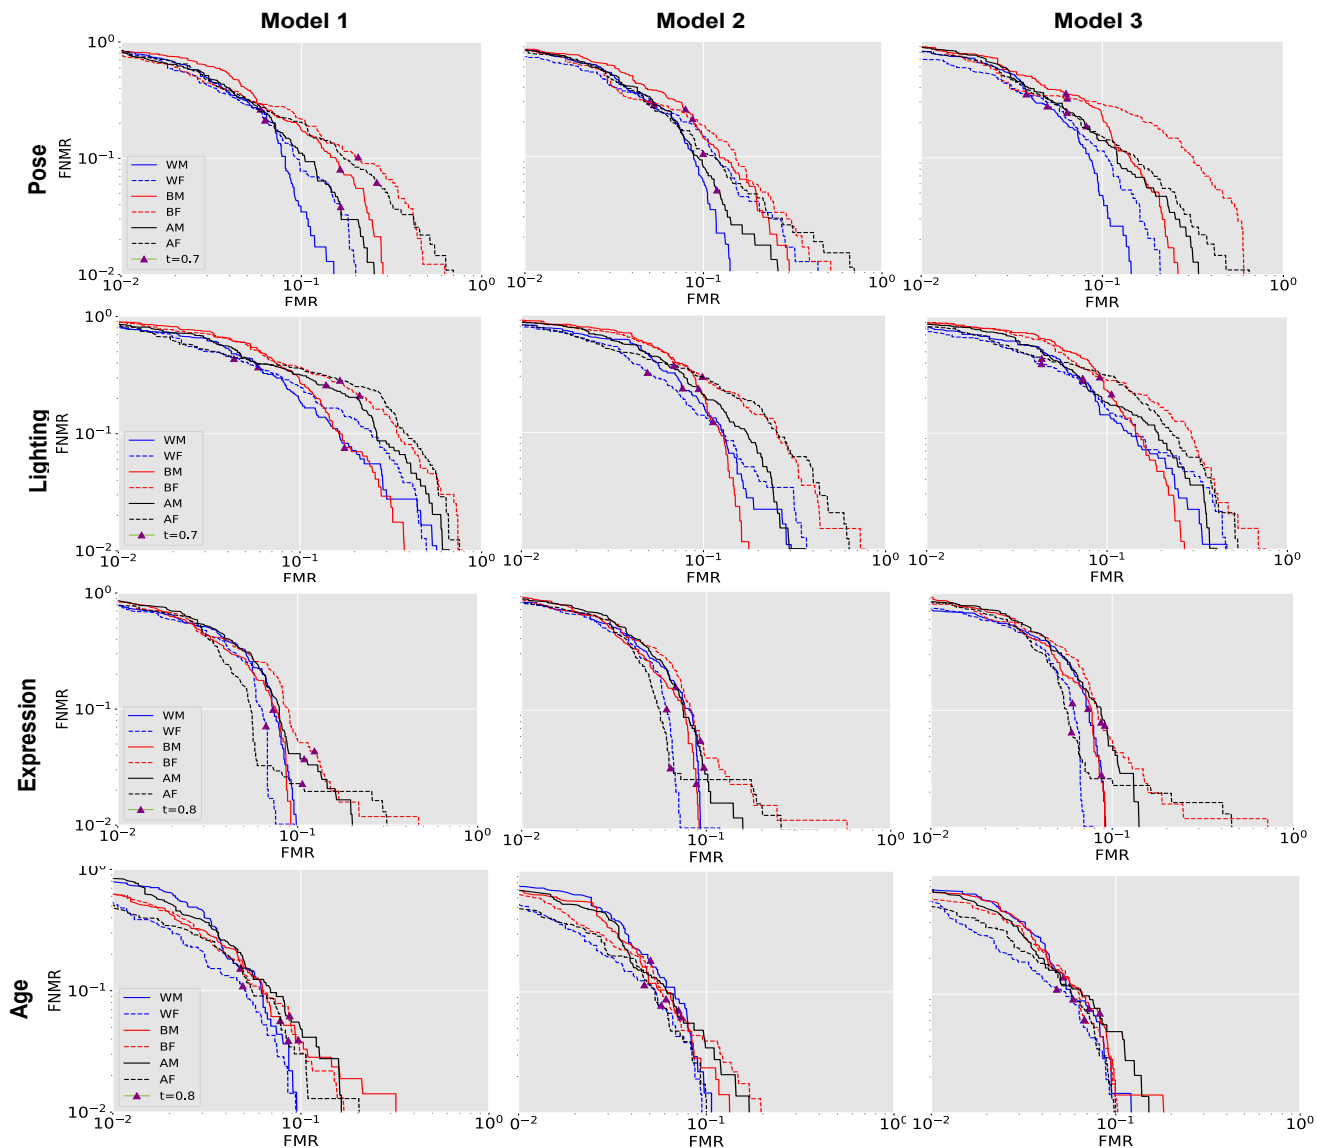


Figure 6. **FNMR vs. FMR** plots from all image pairs using HCIC with $t_{hcic} = 0.3$ as the ground truth labels (see Sec. 4.1.1). Each row shows FNMR vs. FMR vis-a-vis changes in one attribute: pose, lighting expression and age. Each column refers to one of three public models (refer to Sec. 4 for the models’ details). See Sec. 4.1.1 for analysis of these results. Similar to [17], we use purple triangles to show how a *fixed* threshold t (see caption of Fig. 5 for details of t) can affect the performances of different demographic groups in different ways. Different choices of t_{hcic} will not affect the overall trend (see Fig. 9, 10 in Supplementary for results using different t_{hcic} values).

allowing us to remove any unrealistic or artifact-heavy images. More than 70% of images receive a score of “Likely real” or “real for sure” across all demographic groups. See Fig. 3 in Supplementary for details.

Next, we evaluated the humans’ identity distance annotations with respect to each unprotected attribute. Fig. 4 shows distributions of the annotations for pose (plots of remaining attributes are in Fig. 8 in Supplementary), broken down by race/gender subgroups and positive (a) and negative (b) face pairs. The annotators generally agreed with our intended synthesis, but there are also a significant number

of images where they did not. These may be caused either by errors in synthesis or natural human perceptual variability. We show examples of image pairs for which annotators disagreed with our intention in Fig. 11 in Supplementary. Human annotator identity evaluations per image were also fairly consistent. See Fig. 7 in Supplementary for a plot of the median standard deviations and 15% – 85% confidence interval for all attributes. All medians are roughly 0.3 (scale of $[0, 1]$), indicating good consistency among annotators.

4.1.1 Bias analysis of algorithms

Next, we report our main experimental results: bias analysis of the face recognition algorithms. We only use synthetic face pairs where both faces receive an uncanniness score below 0.8, resulting in 11,682 positive face pairs and 35,406 negative face pairs. Please refer to Fig. 12 in Supplementary for examples of face images which receive a score higher than 0.8.

We first evaluate the effect of changes to each attribute on each model’s identity similarity prediction. Fig. 5 presents our results for the pose attribute, on the ResNet34-MS1MV3 model (see Fig. 4, 5, 6 in Supplementary for plots of other attributes and models). As expected, the model reports highest similarity between faces from the same prototype (same demographic group, same seed), and least similarity for faces from different groups and seeds. In addition, we see a decline in reported similarity as the pose angle moves away from 0° , indicating that pose has a causal effect on the similarity score. This result also confirms another hypothesis: by comparing the second and third figure in Fig. 5, it is straightforward to draw the conclusion that the model uses demographic group as essential information when making predictions. Other attributes also have similar causal effects on similarity score, and these effects are also consistent across test on different face recognition models, please see Fig. 5, 6 in Supplementary for details.

We also report False Non Match Rate (FNMR) vs. False Match Rate (FMR) for all models with respect to changes in their decision threshold in Fig. 6, broken down by unprotected attributes. FNMR and FMR are calculated as: $\text{FNMR} = \frac{|\text{false reject pairs}|}{|\text{positive pairs}|}$ and $\text{FMR} = \frac{|\text{false match pairs}|}{|\text{negative pairs}|}$. As we vary the decision threshold t of the face recognition model, we get a different number of false reject/false match pairs, yielding the curves in the figure.

To decide whether a given pair of images belongs to the positive or negative class, we estimate the ground truth label from the human annotation scores. First, we compute the HCIC for each image pair by removing the top and bottom two scores from the 9 human annotations, and compute the average of the remaining five. Next, we choose a threshold t_{hcic} and assign a positive label to pairs for which $\text{HCIC} \leq t_{hcic}$, and negative otherwise.

Fig. 6 shows that the models have clear biases. All models have lowest errors for the White Male and White Female groups. Results on pose changes indicate that all 3 models have significant biases towards race, with Model 3 performing significantly worse on Black Females. For lighting changes, model performances are reasonably consistent across demographic groups, though Model 2 clearly performs worse on Black Female and Asian Female. For expression, all three models perform best on White Males. For age, Model 1 is again significantly better on the White

group, and all three models perform the worst on Black Males. Model biases with respect to gender are less consistent, though clearly present for some scenarios.

5. Discussion and Conclusions

We presented the first experimental approach to measure face recognition algorithmic bias based on generating synthetic images with attributes modified independently, and an identity ground truth computed from human annotator consensus. Our synthetic test dataset is obtained by generating 100 six-face balanced “prototypes” expressing two genders and three races, and systematically varying unprotected attributes (pose, lighting, expression and age). Our final synthetic dataset is extensive, consisting of 12,000/36,000 positive/negative face pairs, and over 500,000 human annotations for both single image attributes and pairwise identity comparisons.

We validated our method by assessing three popular public face recognition models for bias. A first finding is that race and gender deeply affect recognition algorithms: it is much more difficult to recognize identity in tests within the same race and gender groups, then in test where faces belong to different groups. A second finding is that all models are biased: accuracy is better for White Males and Females and (in Model 3) lower for Black Females. Another observation is that face recognition algorithms are more affected by variations in pose and expression, and less so by age and lighting.

Human annotations are consistent and reproducible (See Fig. 7 in Supplementary). They confirm our attribute manipulations (single attribute annotations (Fig. 3)), and show that perceived identity is preserved for same-prototype manipulations (identity comparisons between face pairs (Fig. 4)). Furthermore, the accuracy of these annotators are consistent across race and gender subgroups (See Fig. 8 in Supplementary). This suggests that human perception can provide a robust and reproducible ground truth with which we may benchmark algorithms.

Our work use human annotations to estimate face identity, and the annotation process can be improved. First, we could not ensure that our annotators come from different regions of the world. It is well known that humans are better able to identify others within their own demographic group than outside of it. While we do not see demographic biases in our annotation histograms, it would be good to make sure that annotators come from diverse demographics. Furthermore, it will be interesting to measure annotator response to same/different faces in images of real people and compare these responses to our synthetic image measurements. It is also possible that annotators may latch onto image factors that we did not foresee to make their judgments, such as when two faces generated in the same latent traversal have near-identical hairstyles. Hairstyle does affect face recog-

dition in most models, but SOTA models may be able to ignore hair. We have explored this by using BiSeNet segmentation [51] to remove hair and obtained similar trends to our current results.

Our generation process also has limitations. While the method based on latent space disentanglement has provided a simple and effective way to modify attributes, it can also generate artifacts or change other attributes in addition to the one(s) we intended (See Fig. 12 in Supplementary). This is exacerbated when trying to modify multiple attributes at a time [6]. A valuable future direction would be to thoroughly compare different state-of-the-art generative models (GANs and diffusion models in particular) in terms of their ability to precisely control one or more face attributes.

References

- [1] Vitor Albiero, Kai Zhang, and Kevin W Bowyer. How does gender balance in training data affect face recognition accuracy? In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2020. 2
- [2] Vitor Albiero, Kai Zhang, Michael C King, and Kevin W Bowyer. Gendered differences in face recognition accuracy explained by hairstyles, makeup, and facial morphology. *IEEE Transactions on Information Forensics and Security*, 17:127–137, 2021. 2
- [3] Xiang An, Jiangkang Deng, Jia Guo, Ziyong Feng, Xuhan Zhu, Yang Jing, and Liu Tongliang. Killing two birds with one stone: Efficient and robust training of face recognition cnns by partial fc. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 5
- [4] Xiang An, Xuhan Zhu, Yuan Gao, Yang Xiao, Yongle Zhao, Ziyong Feng, Lan Wu, Bin Qin, Ming Zhang, Debing Zhang, et al. Partial fc: Training 10 million identities on a single machine. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1445–1449, 2021. 2
- [5] Guha Balakrishnan, Raghudeep Gadde, Aleix Martinez, and Pietro Perona. Rayleigh eigendirections (reds): Nonlinear gan latent space traversals for multidimensional features. In *European Conference on Computer Vision*, pages 510–526. Springer, 2022. 3
- [6] Guha Balakrishnan, Yuanjun Xiong, Wei Xia, and Pietro Perona. Towards causal benchmarking of bias in face analysis algorithms. In *Deep Learning-Based Face Analytics*, pages 327–359. Springer, 2021. 2, 4, 9
- [7] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018. 2
- [8] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 67–74. IEEE, 2018. 5
- [9] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. 2, 4
- [10] Patrick Chiroro and Tim Valentine. An investigation of the contact hypothesis of the own-race bias in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, 48(4):879–894, 1995. 1
- [11] Sam Corbett-Davies and Sharad Goel. The measure and mis-measure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*, 2018. 2
- [12] Tiago de Freitas Pereira and Sébastien Marcel. Fairness in biometrics: a figure of merit to assess biometric verification systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(1):19–29, 2021. 2
- [13] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. 1, 2, 5
- [14] Jiankang Deng, Jia Guo, Debing Zhang, Yafeng Deng, Xiangju Lu, and Song Shi. Lightweight face recognition challenge. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2, 5
- [15] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. Fairness in deep learning: A computational perspective. *IEEE Intelligent Systems*, 36(4):25–34, 2020. 2
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2
- [17] Patrick Grother, Mei Ngan, and Kayee Hanaoka. *Face recognition vendor test (fvrt): Part 3, demographic effects*. National Institute of Standards and Technology Gaithersburg, MD, 2019. 1, 2, 7
- [18] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016. 2
- [19] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems*, 33:9841–9850, 2020. 3, 4
- [20] Peter J Hills and J Michael Pake. Eye-tracking the own-race bias in face recognition: Revealing the perceptual and socio-cognitive mechanisms. *Cognition*, 129(3):586–597, 2013. 1
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [22] Kareem J Johnson and Barbara L Fredrickson. “we all look the same to me” positive emotions eliminate the own-race bias in face recognition. *Psychological science*, 16(11):875–881, 2005. 1
- [23] Kimmo Karkkainen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1548–1558, 2021. 2, 4

- [24] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2
- [25] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. 2
- [26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 2
- [27] Yury Kartynnik, Artsiom Ablavatski, Ivan Grishchenko, and Matthias Grundmann. Real-time facial surface geometry from monocular video on mobile gpus. *arXiv preprint arXiv:1907.06724*, 2019. 4
- [28] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*, pages 2564–2572. PMLR, 2018. 2
- [29] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [30] KS Krishnapriya, Vitor Albiero, Kushal Vangara, Michael C King, and Kevin W Bowyer. Issues related to face recognition accuracy varying based on race and skin tone. *IEEE Transactions on Technology and Society*, 1(1):8–20, 2020. 2
- [31] Hao Liang, Lulan Yu, Guikang Xu, Bhiksha Raj, and Rita Singh. Controlled autoencoders to generate faces from voices. In *International Symposium on Visual Computing*, pages 476–487. Springer, 2020. 2
- [32] Zinan Lin, Hao Liang, Giulia Fanti, Vyas Sekar, Rahul Anand Sharma, Elahe Soltanaghaei, Anthony Rowe, Hun Namkung, Zaoxing Liu, Daehyeok Kim, et al. Raregan: Generating samples for rare classes. *arXiv preprint arXiv:2203.10674*, 2022. 2
- [33] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 1, 2, 5
- [34] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 2
- [35] Boyu Lu, Jun-Cheng Chen, Carlos D Castillo, and Rama Chellappa. An experimental evaluation of covariates effects on unconstrained face verification. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 1(1):42–55, 2019. 2
- [36] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 2
- [37] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015. 2
- [38] P Jonathon Phillips, Amy N Yates, Ying Hu, Carina A Hahn, Eilidh Noyes, Kelsey Jackson, Jacqueline G Cavazos, Géraldine Jeckeln, Rajeev Ranjan, Swami Sankaranarayanan, et al. Face recognition accuracy of forensic examiners, superrecognizers, and face recognition algorithms. *Proceedings of the National Academy of Sciences*, 115(24):6171–6176, 2018. 1
- [39] Jean Ponce, Tamara L Berg, Mark Everingham, David A Forsyth, Martial Hebert, Svetlana Lazebnik, Marcin Marszałek, Cordelia Schmid, Bryan C Russell, Antonio Torralba, et al. Dataset issues in object recognition. In *Toward category-level object recognition*, pages 29–48. Springer, 2006. 2
- [40] Novi Quadrianto, Viktoriia Sharmanska, and Oliver Thomas. Discovering fair representations in the data domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8227–8236, 2019. 2
- [41] Vikram V Ramaswamy, Sunnie SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9301–9310, 2021. 2
- [42] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2287–2296, 2021. 3, 4
- [43] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 1, 2
- [44] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020. 3, 4
- [45] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528. IEEE, 2011. 2
- [46] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2
- [47] Sriram Vasudevan and Krishnamurthy Kenthapadi. Lift: A scalable framework for measuring fairness in ml applications. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2773–2780, 2020. 2
- [48] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 692–702, 2019. 2
- [49] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5310–5319, 2019. 2
- [50] Yandong Wen, Weiyang Liu, Adrian Weller, Bhiksha Raj, and Rita Singh. Sphreface2: Binary classification is

- all you need for deep face recognition. *arXiv preprint arXiv:2108.01513*, 2021. [2](#)
- [51] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018. [4](#), [9](#)
- [52] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International conference on machine learning*, pages 325–333. PMLR, 2013. [2](#)
- [53] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340, 2018. [2](#)
- [54] Wenyi Zhao, Rama Chellappa, P Jonathon Phillips, and Azriel Rosenfeld. Face recognition: A literature survey. *ACM computing surveys (CSUR)*, 35(4):399–458, 2003. [1](#), [2](#)
- [55] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David W Jacobs. Deep single-image portrait relighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7194–7202, 2019. [4](#)
- [56] Dominik Zietlow, Michael Lohaus, Guha Balakrishnan, Matthäus Kleindessner, Francesco Locatello, Bernhard Schölkopf, and Chris Russell. Leveling down in computer vision: Pareto inefficiencies in fair deep classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10410–10421, 2022. [2](#)
- [57] James Zou and Londa Schiebinger. Ai can be sexist and racist—it’s time to make it fair, 2018. [2](#)